The impact of centrality on cooperative processes

Sandro M. Reia,[1] Sebastian Herrmann,[2] and José F. Fontanari[1]

[1] *Instituto de Física de São Carlos, Universidade de São Paulo, Caixa Postal 369, 13560-970 São Carlos, São Paulo, Brazil*

[2] *Department of Information Systems and Business Administration, Johannes Gutenberg-Universität, Jakob Welder-Weg 9, 55128 Mainz, Germany*

The solution of today's complex problems requires the grouping of task forces whose members are usually connected remotely over long physical distances and different time zones, so the importance of understanding the effects of imposed communication patterns (i.e., who can communicate with whom) on group performance. Here we use an agent-based model to explore the influence of the betweenness centrality of the nodes on the time the group requires to find the global maxima of families of NK-fitness landscapes. The agents cooperate by broadcasting messages informing on their fitness to their neighbors and use this information to copy the more successful agent in their neighborhood. We find that for easy tasks (smooth landscapes) the topology of the communication network has no effect on the performance of the group and that the more central nodes are the most likely to find the global maximum first. For difficult tasks (rugged landscapes), however, we find a positive correlation between the variance of the betweenness among the network nodes and the group performance. For these tasks the performances of individual nodes are strongly influenced by the agents dispositions to cooperate and by the particular realizations of the rugged landscapes.

arXiv:1610.09939v1 [physics.soc-ph] 31 Oct 2016

I. INTRODUCTION

Problem solving by task-groups represents a substantial portion of the economy of developed countries nowadays [1]. Among the work relationship issues that emerge in this situation, the most important is perhaps that of intra-group communication. In fact, the question "What effect do communication patterns have upon the operation of groups?" has prompted a series of experimental studies in the 1950s, which produced somewhat conflicting conclusions [2–7]. Of particular interest is the case of imposed communication patterns, which happens in the military and industrial organizations, for instance, and in which the researchers determine who can communicate with whom, thus excluding a priori the alternative of self-organization of the group members.

Here we address the issue of the influence of a fixed communication pattern on group performance, as measured by the time the group needs to find the solution of a task. Already in the pioneer studies of the 1950s the concept of centrality has emerged as the chief (but not the sole) determinant of the differences in performance of the various group organizations [2, 3]. Centrality or, more precisely, betweenness centrality is a measure of the closeness of a member to all other group members and, hence, is a measure of the availability of the information necessary for solving the task [8]. In fact, a typical finding of those studies was that the most central position in a pattern (e.g., in a wheel), which is closest to all other positions, is most likely to hit the solution first [3].

Rather than studying small groups of human subjects as in those seminal works, here we consider agent-based simulations aiming at offering a more complete understanding of the interplay between the centrality of the communication patterns and the complexity of the task. Even though it is debatable that conclusions drawn from such an approach may apply to groups of human workers (see, e.g., [9]), they certainly hold for distributed computational systems that are ubiquitous in today's society [10–12].

In particular, we consider a distributed cooperative problem solving model in which agents cooperate by broadcasting messages informing on their partial success towards the completion of the goal and use this information to copy parts of the tentative answer exhibited by the more successful agent in their influence networks [13]. Since copying is an essential ingredient of social learning (i.e., learning through observation) and is central to the remarkable success of our species [14, 15], we expect that our conclusions may be of relevance to the organization of real-world task-groups. The parameters of the model are the number of agents in the system L and the copy propensity $p \in [0, 1]$ that is the same for all agents.

The relevant network metric to our study is the betweenness centrality, which measures a node's centrality in a communication pattern [8]. In order to single out its influence on the group performance, we follow [9] and fix the group size to $L = 16$ and the degrees of the nodes to $k = 3$ (see Fig. 1). In the rest of the paper we will use the terms communication pattern and network interchangeably. The task posed to the agents is to find the unique global maximum of a fitness landscape, whose state space is much larger than the group size L. The difficulty of the task is gauged by the number and distribution of local maxima in the landscape.

We find that for easy tasks (i.e., for landscapes without local maxima) the network topology has no apparent effect on the group performance, but for difficult tasks the strength of the performance is associated to the variance of the betweenness centrality among the network nodes. The network that maximizes that variance exhibits a hierarchical organization with a central node and a modular structure (network A in Fig. 1). It is interesting that such an organization performs better than a more

equalitarian pattern in which the betweenness centrality of all nodes is maximized (network B in Fig. 1). Moreover, we find that the best performances are achieved by the so-called inefficient networks, which are characterized by long average path lengths that delay the propagation of information through the network. This is because in rugged landscapes the information on fitness exchanged by the agents is often misleading, hinting to the locations of local maxima, rather than to the position of the global maximum of the fitness landscape.

In addition, a more detailed consideration of the performance of the nodes shows that for easy tasks the chance that a node finds the answer first is positively correlated with its betweenness centrality. For difficult tasks, however, the chance of a node hitting the solution depends on the copy propensity p of the agents: for small p all agents are roughly equiprobable of finding the solution. Near the value of p that optimizes the group performance, the central agents perform better. For large p the more peripheral nodes have a better chance to get the answer first.

The rest of this paper is organized as follows. In Section II we offer an outline of the NK model of rugged fitness landscapes [16] which we use to represent the tasks presented to the agents. The behavioral rules that guide the agents in their searches for the global maximum of the landscapes are explained in Section III and the four fixed communication patterns the agents use to exchange information on their tentative solutions are introduced in Section IV. In Section V we present and analyze the results of our simulations, emphasizing the comparative performance between the different patterns. Finally, Section VI is reserved to our concluding remarks.

II. TASK

The task posed to a system of L agents $i = 1, \ldots, L$ is to find the unique global maximum of a fitness landscape generated using the NK model [16]. The NK model is the paradigm for problem representation in organizational theory [17–20], since it allows the tuning of the ruggedness of the landscape – and hence of the difficulty of the task – by changing the integer parameters N and K. More pointedly, the NK landscape is defined in the space of binary strings of length N and so this parameter determines the size of the state space, namely, 2^N. The other parameter $K = 0, \ldots, N-1$ is the degree of epistasis that influences the ruggedness of the landscape for fixed N. For $K = 0$ the (smooth) landscape has a single maximum, whereas for $K = N-1$, the (uncorrelated) landscape has on the average $2^N/(N+1)$ maxima with respect to single bit flips and the NK model reduces to the Random Energy model [21, 22]. Finding the global maximum of the NK model for $K > 0$ is an NP-complete problem [23], which means that the time required to solve *all* realizations of that landscape using any currently known deterministic algorithm increases exponentially fast with the length N of the strings [24]. We refer the reader to the original paper by Kauffman and Levin for details on the procedure to generate a random realization of an NK landscape [16].

TABLE I. Statistics of the number of maxima in the sample of 100 NK-fitness landscapes with $N = 16$ used in the computational experiments.

K	mean	min	max
0	1	1	1
3	84.24	43	132
7	664.51	573	770

Since our goal is to compare the performances of cooperative problem-solving systems using the four communication patterns shown in Fig. 1, we must guarantee that they search the same realizations of the fitness landscapes, as distinct landscape realizations may differ greatly in the number of local maxima for $K > 0$. Thus for each set of the parameters N and K we generate and store 100 landscape realizations, which we use to test the four patterns. In particular, we fix the string length to $N = 16$ and allow the degree of epistasis to take on the values $K = 0$, 3, and 7. Table I shows the mean number of maxima for each sample of 100 landscapes as well as two extreme values, namely, the minimum and the maximum number of maxima in the sample.

III. COOPERATIVE SEARCH

Once the task is specified we can set up the rules for the operation of the agents. We assume that initially each agent operates in a binary string drawn at random with equal probability for the bits 0 and 1. The search begins with the selection of an agent at random, say agent i, at time $t = 0$. This agent can choose between two distinct processes to operate on its string, as described next [13, 25].

The first process, which happens with probability $1-p$, is the elementary move in the solution space that consists of picking a bit at random from the string and flipping it. This elementary move allows the agents to explore in an incremental way the 2^N-dimensional state space. The second process, which happens with probability p, is the imitation of a model string, which is defined as the string that exhibits the largest fitness among the (fixed) subgroup of agents that can influence (i.e., are connected to) agent i. The model string and the string i (i.e., the string operated by agent i) are compared and the different bits are singled out. Then agent i selects at random one of the distinct bits and flips it so that this bit is now the same in both strings. In the case the string i is identical to the model string, agent i executes the elementary move (flipping a random bit) with probability one.

After agent i has updated its string, we increment the time t by the quantity $\Delta t = 1/L$. Since a string operation

always results in a change of fitness, we need to recalculate the model string for all the subnetworks that contain agent i. Then another agent is selected at random and the procedure described above is repeated. Note that during the increment from t to $t+1$, exactly L string operations are performed, though not necessarily by L distinct agents.

The search ends when one of the agents finds the global maximum and we denote by t^* the halting time. The efficiency of the search is measured by the total number of string operations necessary to find that maximum, i.e., Lt^* [3] (see also [11, 12]) and so the computational cost of a search is defined as $C \equiv Lt^*/2^N$, where for convenience we have rescaled t^* by the size of the solution space 2^N.

The parameter $p \in [0, 1]$ is the copy propensity of the agents. The case $p = 0$ corresponds to the baseline situation in which the agents explore the state space independently of each other. The copy or imitation procedure described above was based on the incremental assimilation mechanism used to study the influence of external media [26, 27] in Axelrod's model of social influence [28]. Here we assume that the L agents are identical with respect to their copy propensities (see [29] for the relaxation of this assumption).

A word is in order about a similar agent-based model used in organizational theory [18, 19]. In that model the agents always copy the fittest string in their neighborhood (i.e., $p = 1$, but see below), a move called exploitation, except when the agent is the fittest one, in which case a single bit is flipped randomly but, differently from our elementary move, this change is enacted only if it increases the fitness of the agent. This move is called exploration. In addition, in the case the agent imitates a more successful neighbor, it copies the entire string by changing many bits simultaneously. This is then a non-incremental move in the state space. As a result, the search may permanently get stuck in a local maximum of the landscape [18, 19]. This the reason the performance measure in those studies is taken as the average fitness of the group after a fixed search time or as the fraction of searches that found the global maximum for unlimited search times.

IV. COMMUNICATION PATTERNS

Here we focus on a commonly studied network metric, namely, the betweenness centrality which is a measure of a node's centrality in a communication pattern or network [8]. More pointedly, the betweenness centrality of node i, which we denote by $\mathcal{B}_i$, is given by the ratio between the number of shortest paths from all nodes to all others that pass through node i and the number of shortest paths from all nodes to all others regardless of whether they pass through node i or not. Clearly, a node with high betweenness centrality has a large influence on the transfer of information through the network. Hence it is relevant to understand the role of this metric on the performance of distributed cooperative problem-solving systems. We can introduce global metrics as well, such as the average betweenness centrality $\mu_{\mathcal{B}} = \sum_{i=1}^{L} \mathcal{B}_i/L$ where L is the number of nodes of the network. The variance $\sigma^2_{\mathcal{B}}$ of the betweenness centrality is similarly defined. For simplicity, henceforth we will refer to the betweenness centrality as simply the betweenness, since the other type of betweenness, namely, edge betweenness [30] will not be considered here.

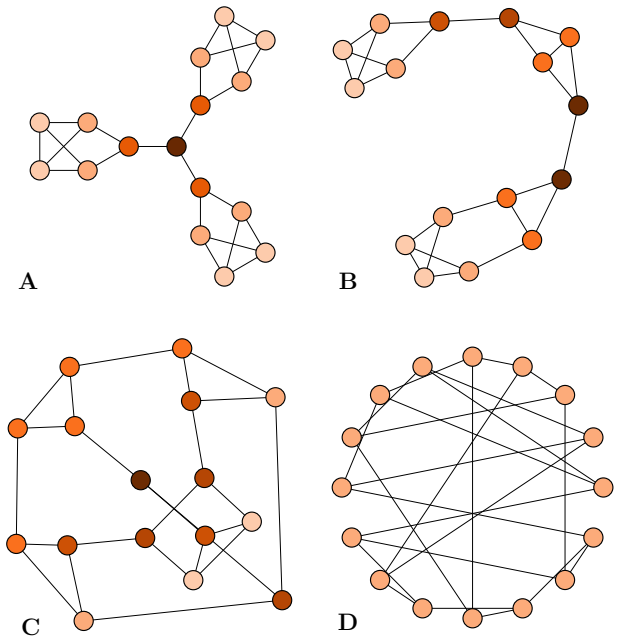

FIG. 1. (Color online) The four network topologies with $L = 16$ nodes and fixed degree $k = 3$ used in the computational experiments. The darker the shade of a node, the higher its betweenness. Network A is the topology that maximizes both the maximum betweenness and the betweenness variance, network B maximizes the average betweenness, network C is a typical random network regarding the average betweenness, and network D minimizes both the maximum betweenness and the betweenness variance (see Table II).

For networks with $L = 16$ nodes and fixed degree $k = 3$ we can obtain, through an exhaustive search on the space of networks, three networks with special properties regarding the betweenness metric, which are exhibited in Fig. 1. For instance, network A is the network that maximizes the maximum betweenness among the L nodes. It happens also that this network exhibits the maximum variance (or standard deviation) of the betweenness among the nodes. Network B exhibits the maximum average betweenness. Network D minimizes the maximum betweenness among the L nodes and exhibits also the minimum variance of the betweenness among the nodes. Actually, all nodes have the same betweenness in this network. Network C is a typical random network, regarding the average betweenness, with $L = 16$ and $k = 3$. To obtain this network we generated a sample of 10^5 random

TABLE II. Summary statistics of the networks' betweenness and average path length for the four topologies used in the computational experiments.

Network	$\mu_{\mathcal{B}}$	$\sigma_{\mathcal{B}}$	$\bar{l}$
A	0.1678	0.2118	3.35
B	0.2047	0.1862	3.87
C	0.1036	0.0348	2.45
D	0.0857	0	2.2

networks, calculated the average betweenness of each network, and then the average of the sample: network C was the network whose average betweenness was closest to the sample average. Networks A and B were considered in the study of Ref. [9]. Table II exhibits the average betweenness ($\mu_{\mathcal{B}}$) and the standard deviation ($\sigma_{\mathcal{B}}$) of these four networks. The networks were ordered from high to low values of their betweenness variances.

Another metric of interest used to characterize the communication patterns is the average path length $\bar{l}$ defined as the average number of steps along the shortest paths for all possible pairs of network nodes. Because it is a measure of the efficiency of the flow of information on a network [31], it has been used to classify the communication patterns as efficient (short path lengths) and inefficient (long path lengths) [9]. In that sense, the two networks shown in the upper row of Fig. 1 are classified as inefficient networks and those shown in the lower row as efficient networks (see Table II).

V. RESULTS

For a given communication pattern and for fixed values of the NK-model parameters we proceed as follows. For each realization of a fitness landscape we carry out 10^4 to 10^5 searches starting from different initial conditions (initial strings) and the resulting average computational cost is then averaged again over 100 distinct landscapes. We recall that the four networks of Fig. 1 are tested on the same landscapes. The error bars are smaller than the symbol sizes in all figures shown in this section.

In Fig. 2 we show the dependence of the average computational cost $\langle C\rangle$ on the copy propensity p of the agents for increasing task difficulties as measured by the landscapes' ruggedness, $K = 0$, 3 and 7. For landscapes with no local maxima ($K = 0$, upper panel of Fig. 2), the performances of the four topologies are practically indistinguishable in the scale of the figure and the mean computational cost decreases with increasing p, i.e., the best performance is attained by always copying the model string ($p = 1$) and allowing only its clones to explore the landscape through the elementary move.

The presence of a moderate number of local maxima ($K = 3$, middle panel of Fig. 2) impacts the performance only for large values of the copy propensity and the effect is more pronounced for the more efficient networks C and D. The performances of networks A and B are very similar, except in the region of p very close to 1, where network A slightly outperforms network B. The difference between the performances of these two networks becomes evident for difficult problems ($K = 7$, lower panel of Fig. 2) only, as illustrated in the inset of that panel.

It seems then that the determinant factor for the superior performance of a communication pattern is the variance of the betweenness among the nodes (see Table II). However, an alternative explanation may be the presence of modules in network A and quasi-modular structures in networks B and C but not in network D, which exhibits the worst performance. In fact, it has been argued that the modular organization, which is characteristic of hierarchical networks, may facilitate the escape from the local maxima [32] (see [33] for experimental evidence on the effect of a hierarchical social network structure on the efficiency of collective action). The group performance correlates negatively with the efficiency, as measure by $\bar{l}$, of the networks in the case of rugged landscapes. This happens because the model agents may broadcast misleading information in that case, and so it is advantageous to slow down the information transmission so as to allow the agents more time to explore the solution space away from the neighborhoods of the local maxima.

We find that $\langle C\rangle$ is quite insensitive to variations on the topology of the network for small values of p. In particular, for $p = 0$ one recovers the results of the independent search, $\langle C\rangle \approx 1.08$ (see [25] for an analytical estimate of this value), regardless of the network topology and of the value of the parameter K. As the difficulty of the task increases the optimum copy propensity decreases towards zero and the minimum of $\langle C\rangle$ becomes shallower. In particular, for $K = N - 1$ that minimum happens at $p = 0$. Finally, we note that since finding the global maxima of NK landscapes with $K > 0$ is an NP-Complete problem [23], one should not expect that the imitative search (or any other search strategy, for that matter) would find those maxima for a large sample of landscapes much more rapidly than the independent search.

Figure 2 reveals the superior performance of networks whose nodes exhibit the largest variability of betweenness in long runs, i.e., when there is no limit to the duration of the search. Now we examine whether this finding holds also in the case where a maximum search time t or computational cost C is fixed a priori. Figure 3 shows the fraction of runs $F(C)$ that found the global maximum for a fixed value of C in difficult tasks, i.e., landscapes with $N = 16$ and $K = 7$. For easy tasks ($K = 0$) we find that, similarly to the situation for long runs, the fixed-cost performances of the four networks are practically indistinguishable (data not shown). Hence, the conclusions for the long runs displayed in Fig. 2 are valid in the case that the computational cost of the search is fixed a priori as well.

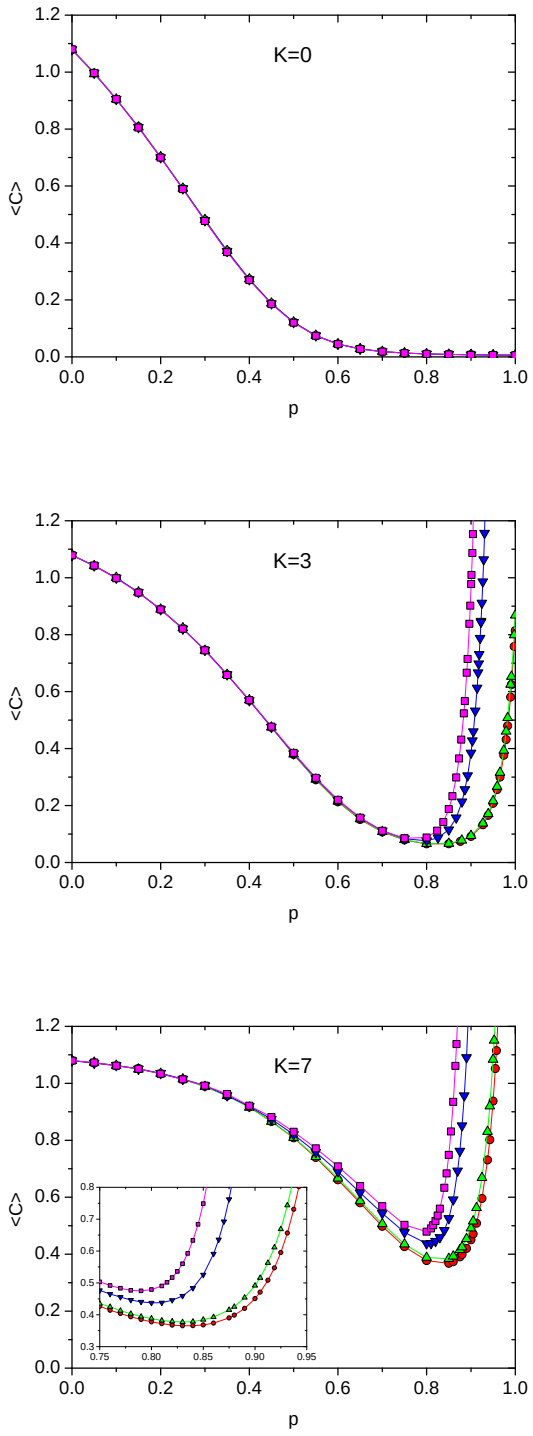

FIG. 2. (Color online) Average computational cost $\langle C \rangle$ as function of the copy propensity p for the four communication patterns shown in Fig. 1 according to the convention: network A (●), network B (▲), network C (▼) and network D (■). Each panel shows the results for a fixed value of the parameter K that measures the problem difficulty (see Table II), from very easy ($K = 0$) to difficult ($K = 7$). The lines are guides to the eye.

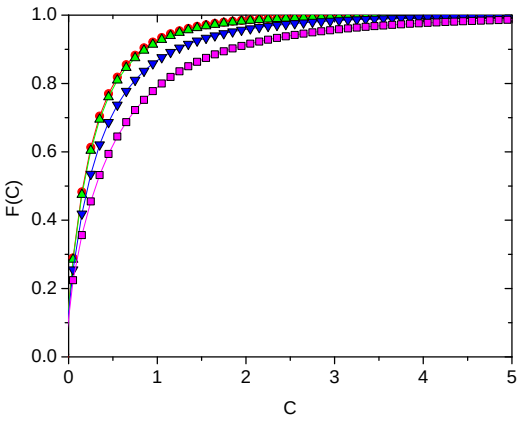

FIG. 3. (Color online) Fraction of the searches that found the solution for a fixed computational cost C for network A (●) network B (▲), network C (▼) and network D (■). The copy propensity is $p = 0.85$ and the parameters of the landscape are $N = 16$ and $K = 7$. The lines are guides to the eye.

For network A, which exhibits nodes with 4 distinct betweenness values, namely, $\mathcal{B} = 0.003175$, 0.114286, 0.422222, and 0.714286 with degeneracies $D_{\mathcal{B}} = 6$, 6, 3 and 1, respectively, we can measure the chance $F_{\mathcal{B}}(C)$ that a specific node with a certain betweenness finds the global maximum for a fixed computational cost C. This quantity is given by the fraction of runs with computational cost less than C for which a particular node with betweenness $\mathcal{B}$ finds the solution. Since for $C \to \infty$, we can guarantee that all runs have halted, we have $\sum_{\mathcal{B}} D_{\mathcal{B}} F_{\mathcal{B}}(C) = 1$ in this limit. The results for $K = 0$ and $K = 3$ are shown in Figs. 4 and 5, respectively, with the copy propensity set to its maximum value, $p = 1$.

For easy tasks ($K = 0$), the central node of network A is most likely to get the answer first, regardless of the allotted search time as shown in Fig. 4. More generally, the chance of a node hitting the solution increases steadily with its betweenness. This conclusion holds for networks B and C as well and for all values of the copy propensity p. We recall that all nodes of network D are identical regarding their betweenness values, so this network is unfit for this type of analysis. Surprisingly, the situation is reversed for more difficult tasks ($K = 3$) as shown in Fig. 5. Although it is still true that the central node is the most likely to find the solution for short runs (low computational cost), it is the least likely for long runs. In fact, for long runs the lesser the centrality of a node, the greater its chance of hitting the solution. We note, however, the superior performance of the more central nodes for short runs seems to be a peculiarity of network A since for networks B and C the nodes with the lowest betweenness are the most likely to find the solution regardless of the value of C (data not shown). It is interesting to mention a related finding within the context

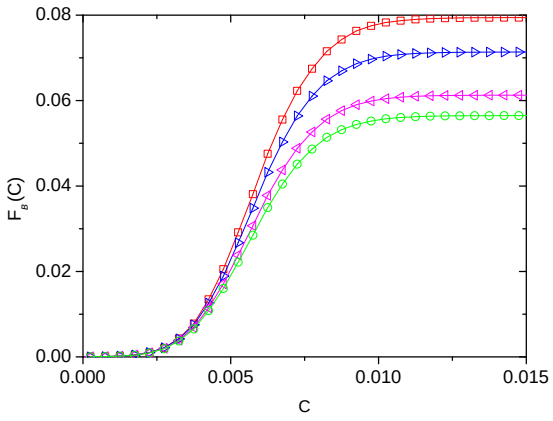

FIG. 4. (Color online) Probability that a node with betweenness $\mathcal{B}$ in network A gets the answer first for runs with computational cost less than C. The convention is $\mathcal{B} = 0.003175(\circ)$, $0.114286(\triangleleft)$, $0.422222(\triangleright)$, and $0.714286(\square)$. The copy propensity is $p = 1$ and the parameters of the smooth landscape are $N = 16$ and $K = 0$. The lines are guides to the eye.

of the spreading of epidemics in complex networks that in certain circumstances the most efficient spreaders may not be the most central individuals in the network [34].

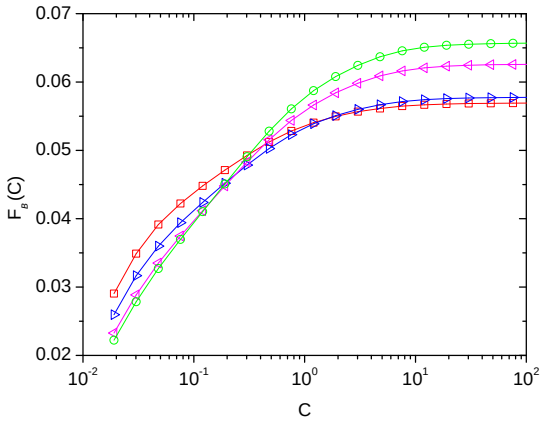

FIG. 5. (Color online) Same as Fig. 4 but for rugged landscapes with parameters $N = 16$ and $K = 3$. The logarithmic scale in the x-axis highlights the change of the networks performance hierarchy for long and short runs.

Actually, for rugged landscapes the performance of a particular node is way more complicated than for smooth landscapes, because it depends on the copy propensity p as shown in Fig. 6. For instance, for p close to the value that minimizes the computational cost (see middle panel of Fig. 2) the central nodes perform better but for p close to 1, in the region where the search is hampered by the local maxima, the peripheral nodes perform better, in agreement with Fig. 5. We note that the highly non-monotonic behavior of the probability $F_{\mathcal{B}}$ for the central nodes is an artifact of averaging over different landscapes realizations. In fact, for some landscape realizations we find a monotonous decrease of $F_{\mathcal{B}}$ with increasing p and for others we find that this probability increases with p so the central nodes end up performing better than the peripheral ones.

Since even our simple agent-based model yields plenty of discordant results regarding the performance of individual nodes on difficult tasks, it is no wonder the profusion of conflicting conclusions drawn from the studies with human subjects, as neither the cooperation strategies used by the subjects nor the difficulty of the problems are controlled variables in those experiments.

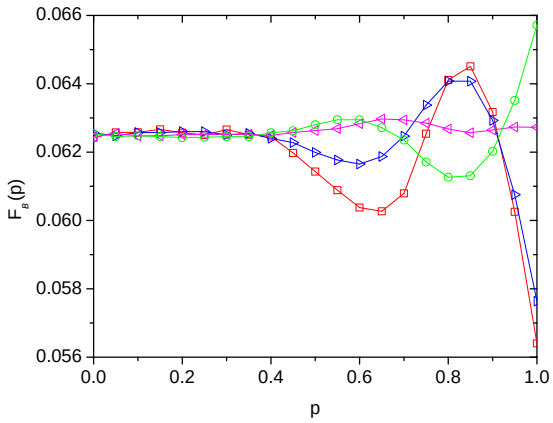

FIG. 6. (Color online) Probability that a node with betweenness $\mathcal{B}$ in network A gets the answer first for time-unrestricted runs as function of the copy propensity. The parameters of the rugged landscapes are $N = 16$ and $K = 3$. The symbols convention is the same as for Fig. 4 and the lines are guides to the eye.

VI. DISCUSSION

The claim that restrictions on the communication channels available to a group affects its problem-solving efficiency is hardly controversial. However, the issue whether there is a communication pattern which gives significantly better performance than others in solving specific or general tasks has produced conflicting findings [2–7, 9, 18, 19]. The main reason seems to be the strong dependence of the group performance on several aspects of the group organization as well as on the complexity of the tasks and on the cooperation strategies used by the subjects. Hence a thorough analysis of the vast parameter space of the group/task composite is necessary to elucidate that issue. Such comprehensive study is feasible only through computational experiments [18].

Here we attempt to clarify the role of the betweenness centrality on the performance of a group as well as on the performance of its individual members. To achieve that we compare networks with the same number of nodes ($L = 16$), which are identical with respect to their degrees ($k = 3$ for all nodes), as illustrated in Fig. 1. The networks were generated so as to exhibit special properties regarding their local and global betweenness metrics [9]. The task posed to the agents is to find the unique global maximum of a NK-fitness landscape with the parameter $N = 16$ fixed but with variable K. The difficulty of the task increases with K due to the proliferation of local maxima. The choice $N = L = 16$ ensures that the group size is fixed close to its optimal value for easy tasks [25] and that the size of the state space 2^N is much larger than L, which makes the search for the single global maximum a challenging task.

We find that for simple tasks there is no significant difference in the performances of the different network topologies. However, the performance of the group members, which is measured by the chance they get the answer first, is strongly correlated to the centrality of the node in the network: the more central a node is, the more likely it is to find the solution of the task. These findings are in agreement with the experimental results [3, 6]. We stress, however, that the unresponsiveness of the easy-task performance to changes on the communication patterns is because the number of nodes as well as the density of links are the same for the four networks of Fig. 1. Variation of these parameters has a large effect on the group performance for easy tasks. For instance, for those tasks there is an optimal group size that minimizes the cost of the search and the optimal performance is achieved by fully connected networks [35], in agreement with the experimental finding that groups with more communication channels perform better for simple tasks [4].

For complex tasks, the best performing communication pattern is the pattern that maximizes the variance of the betweenness among the nodes (network A in Fig. 1). This network also exhibits the node with the largest possible betweenness allowed by the constraints $L = 16$ and $k = 3$ for all nodes. The group performance degrades as that variance decreases towards zero, so that the worst performing pattern is that in which all nodes have the same betweenness (network D in Fig. 1). Interestingly, among the topologies displayed in Fig. 1, network D is the topology with the shortest average path length and so it is the most efficient regarding the transmission of information through the network, whereas network A is above network B only in this rank of efficiency (see Table II). Our finding that inefficient networks perform better is justified by noting that speeding up the transmission of information through the network makes sense only if one can guarantee its faithfulness and usefulness, otherwise it may be wiser to slow it down and give more time for the agents to explore different regions of the state space [18, 36, 37]. We note, however, that even in the case the fitness values provide faithful information on the location of the global optimum, viz. for landscapes with $K = 0$, the so-called efficient networks do not perform better than the inefficient ones.

In addition, the attempt to maximize the betweenness of all nodes results in the network with the largest average betweenness (network B in Fig. 1), which exhibits the second best performance. Thus it seems that the key factor to improve group performance is not the maximization of the number of nodes that have large betweenness but rather the assignment of a large variety of betweenness values to the nodes. Hence, within the perspective of the betweenness centrality metric, diversity is crucial to boost group performance.

Regarding the performance of an individual node of the communication network, measured by the probability that it hits the global maximum first, we find a neat positive correlation between the centrality of a node and its performance on easy tasks, regardless of the topology or of the particularities of the agents, such as their copy propensities p. For complex tasks, however, there is no such a general verdict as the performances of the nodes are strongly influenced by the agents dispositions to cooperate and by the specific realizations of the rugged landscapes. Averaging over the landscape realizations yields a complicated dependence on the parameter p. For instance, the central nodes perform better when the network performance is optimal, whereas the peripheral nodes win when the network performance is most heavily harmed by the local maximum traps. This sensitivity on the details of the model and the consequent impossibility of drawing general (i.e, task and subject independent) conclusions is reminiscent of the conflicting outcomes that characterize the experimental literature dealing with the effects of the communication patterns in task-oriented groups [2–7, 9, 36, 37].

A word is in order about the interesting online experiments conducted using Amazon's Mechanical Turk, in which human subjects (players) select points for oil-drilling on a map [9]. The good and the bad oil wells are the maxima and the minima of a rugged landscape and their locations are unknown to the subjects, who, however, are able to see the coordinates of the selected points as well as the earnings of their network neighbors. The goal of each player is to maximize its own earnings by picking the more productive well. In contrast with our findings, the best performing networks in the web-based experiments are the so-called efficient networks, characterized by short average path lengths, such as networks C and D in Fig. 1. A possible reason for this discrepancy is the distinct performance measure used in those experiments, namely, the average earnings of the group rather than the time to find the optimum oil well. In addition, since the game doesn't stop when that optimum is eventually found by a player, truly useful information about the coordinates of the optimum becomes available to the other players, which may then explain the superiority of the topologies with short average path lengths.

Finally, we note that the study of distributed cooperative problem-solving systems diverges from the game theoretical literature on cooperation that followed Robert Axelrod's 1984 seminal book The Evolution of Cooperation [38]. In fact, in the context of cooperative processes, there is no conflict of interests between the agents and the opposite of cooperation is independent work, rather than defection. In addition, in the game theoretical framework it is usually assumed that mutual cooperation is the most rewarding strategy for the group in the long run, whereas here we argue that too much cooperation, which results from a high value of the copy propensity, may lead to disastrous results, akin to the so-called groupthink phenomenon that happens when everyone in a group starts thinking alike [39].

Most studies of the influence of communication patterns on the operation of groups have considered externally imposed patterns that dictate the distribution of the communication channels among the group members (see [5] for an exception), thus precluding the emergence of a self-organized network. This connectivity-driven network approach is fitting to describe situations where the connections among nodes are persistent features, such as in the Internet, but it may not be so suitable to model the more ephemeral work relationships, which may change on a very short time scale. A promising avenue for further theoretical investigation is to consider dynamic unidirectional links, as in activity driven models of varying networks [40], that can be created or destroyed depending on whether a previous copy process resulted in an increase or decrease of the copier fitness. The topology of the resulting time varying self-organized network may shed light on the issue of the emergence of leadership in a task-force [41].

ACKNOWLEDGMENTS

The research of JFF was supported in part by grant 15/21689-2, São Paulo Research Foundation (FAPESP) and by grant 303979/2013-5, Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq). SMR was supported by grant 15/17277-0, São Paulo Research Foundation (FAPESP).

[1] S.E. Page, The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies, Princeton University Press, Princeton, NJ, 2007.
[2] A. Bavelas, Communication Patterns in Task-Oriented Groups, J. Acoustical Soc. Amer. 22 (1950) 725–730.
[3] H.J. Leavitt, Some Effects of Certain Communication Patterns on Group Performance, J. Abnorm. Soc. Psych. 46 (1951) 38–50.
[4] G.A. Heise, G.A. Miller, Problem Solving by Small Groups using Various Communication Nets., J. Abnorm. Soc. Psych. 46 (1951) 327–335.
[5] H. Guetzkow, H.A. Simon, The Impact of Certain Communication Nets upon Organization and Performance in Task-oriented Groups, Manage. Sci. 1 (1955) 233–250.
[6] M.E. Shaw, Some Effects of Unequal Distribution of Information upon Group Performance in Various Communication Nets, J. Abnorm. Soc. Psych. 49 (1954) 547–553.
[7] M. Mulder, Communication Structure, Decision Structure and Group Performance, Sociometry 23 (1960) 1–14.
[8] L.C. Freeman, A set of measures of centrality based on betweenness, Sociometry 40 (1977) 35–41.
[9] W. Mason, D.J. Watts, Collaborative Learning in Networks, Proc. Natl. Acad. Sci. 109 (2012) 764–769.
[10] B.A. Huberman, The performance of cooperative processes, Physica D 42 (1990) 38–47.
[11] S.H. Clearwater, B.A. Huberman, T. Hogg, Cooperative Solution of Constraint Satisfaction Problems, Science 254 (1991) 1181–1183.
[12] S.H. Clearwater, T. Hogg, B.A. Huberman, Cooperative Problem Solving. In: B.A. Huberman (Ed.), Computation: The Micro and the Macro View. World Scientific, Singapore, pp. 33–70, 1992.
[13] J.F. Fontanari, Imitative Learning as a Connector of Collective Brains, PLoS ONE 9 (2014) e110517.
[14] S. Blackmore, The Meme Machine, Oxford University Press, Oxford, 2000.
[15] L. Rendell, R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M.W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, K.N. Laland, Why Copy Others? Insights from the Social Learning Strategies Tournament, Science 328 (2010) 208–213.
[16] S.A. Kauffman, S. Levin, Towards a general theory of adaptive walks on rugged landscapes, J. Theor. Biol. 128 (1987) 11–45.
[17] D.A. Levinthal, Adaptation on Rugged Landscapes, Manag. Sci. 43 (1997) 934–950.
[18] D. Lazer, A. Friedman, The Network Structure of Exploration and Exploitation, Admin. Sci. Quart. 52 (2007) 667–694.
[19] S. Herrmann, J. Grahl, F. Rothlauf, Problem Complexity in Parallel Problem Solving. In: K. Fischbach, M. Grossmann, U.R. Krieger, T. Staake (Eds.), Proceedings of the International Workshop on Modeling, Analysis and Management of Social Networks and their Applications, University of Bamberg Press, Bamberg, pp. 77–83, 2014.
[20] S, Billinger, N. Stieglitz, T.R Schumacher, Search on Rugged Landscapes: An Experimental Study, Organ. Sci. 25 (2013) 93–108.
[21] B. Derrida, Random-energy Model: An Exactly Solvable Model of Disordered Systems, Phys. Rev. B 24 (1981) 2613–2626.
[22] D.B. Saakian, J.F. Fontanari, Evolutionary dynamics on rugged fitness landscapes: exact dynamics and information theoretical aspects, Phys. Rev. E 80 (2009) 041903.
[23] D. Solow, A. Burnetas, M. Tsai, N.S. Greenspan, On the Expected Performance of Systems with Complex Interactions Among Components, Complex Systems 12 (2000) 423–456.

[24] M.R. Garey, D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, Freeman, San Francisco, CA, 1979.
[25] J.F. Fontanari, Exploring NK Fitness Landscapes Using Imitative Learning, Eur. Phys. J. B 88 (2015) 251.
[26] Y. Shibanai, S. Yasuno, I. Ishiguro, Effects of Global Information Feedback on Diversity, J. Conflict Res. 45 (2001) 80–96.
[27] L.R. Peres, J.F. Fontanari, The media effect in Axelrod's model explained, EPL 96 (2011) 38004.
[28] R. Axelrod, The Dissemination of Culture: A Model with Local Convergence and Global Polarization, J. Conflict Res. 41 (1997) 203–226.
[29] J.F. Fontanari, When more of the same is better, EPL 113 (2016) 28009.
[30] M. Girvan, M.E.J. Newman, Community structure in social and biological networks, Proc. Natl. Acad. Sci. USA 99 (2002) 7821–7826.
[31] R. Albert, A.-L. Barabási, Statistical mechanics of complex networks, Rev. Mod. Phys. 74 (2002) 47–97.
[32] S.M. Reia, J.F. Fontanari, Effect of group organization on the performance of cooperative processes, Ecol. Complex. (2016), http://dx.doi.org/10.1016/j.ecocom.2016.09.002
[33] M. Nagy, Z. Akos, D. Biro, T. Vicsek, Hierarchical group dynamics in pigeon flocks, Nature. 464 (2010) 890–894.
[34] M. Kitsak, L.K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H.E. Stanley, H.A. Makse, Identification of influential spreaders in complex networks, Nature Physics 6 (2010) 888–893.
[35] J.F. Fontanari, F.A. Rodrigues, Influence of network topology on cooperative problem-solving systems, Theory Biosci. 135 (2016) 101–110.
[36] W.A. Mason, A. Jones, R.L. Goldstone, Propagation of innovations in networked groups, J. Exp. Psychol. Gen. 137 (2008) 422–433.
[37] C. Fang, J. Lee, M.A. Schilling, Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning, Organ. Sci. 21 (2010) 625–642.
[38] R. Axelrod, The Evolution of Cooperation, Basic Books, New York,1984.
[39] I.L. Janis, Groupthink: psychological studies of policy decisions and fiascoes. Houghton Mifflin, Boston, 1982.
[40] N. Perra, B. Goncalves, R. Pastor-Satorras, A. Vespignani, Activity driven modeling of time varying networks, Sci. Rep. 2 (2012) 469
[41] W.B. Lindquist, I.D. Chase, Data-based analysis of winner-loser models of hierarchy formation in animals, Bull. Math. Biol. 71 (2009) 556–584.